Specification

A WORKFLOW SYSTEM FOR DETECTION AND CLASSIFICATION OF IMAGES SUSPECTED AS PORNOGRAPHIC

This application is a continuation-in-part of Ser. No. 09/823,139 Filed March 29th, 2001 titled Automated Detection of Pornographic Images.

Background of the Invention

Field of the Invention

This invention relates generally to the field of workflow based image analysis and classification and more particularly to a classification of images suspected as pornographic in nature or images suspected as being of a copyright nature.

Description of Prior Art

A variety of methods have been used in an attempt to detect and categorize objectionable images. Pornographic-free web sites, such as sites targeting families and children have been set up for shielding children from viewing objectionable material. Although a particular site may be pornographic free, and considered acceptable for access by children, it is still possible to gain access to an objectionable web site by starting from an acceptable site. Software applications and Internet services such as Net-Nanny and Cyber-Sitter were created and marketed to help parents prevent their children from accessing objectionable documents by blocking access to specific web sites.

One type of protective software is designed to store the addresses of objectionable web sites, and block access to these sites. Example of prior art are US 5,678,041 to Baker and Grosse, US 6,049,821 to Theriault et. al., and US 6,065,055 to Hughes and Elswick.

Another form of software protection screens the text information accessed by a computer from the network and blocks information sources that are considered objectionable. Examples of such prior art include US 5,832,212 to Cragun & Day, US 5,835,722 to Bradshaw and Shih, US 5,996,011 to Humes, US 6,065,056 to Bradshaw and Shih, and US 6,266,664 to Russell-Falla & Hanson.

Such methods are prone to error as many words have subtle double-meanings which can easily be misinterpreted by such software and other words commonly used in everyday conversation can be easily taken out of context. Further, although such software does have a

role to play in content management it does not address the fundamental issue of determining the nature of graphical content on large image collections such as Internet photo communities.

Yet another type of protective software blocks access to URLs except those that are members of a list of manually approved URLs. Examples of prior art include US 5,784, 564 to Camaisa et. al. and US 6,286,001 to Walker & Webb.

These approaches are not highly effective because it is a practical impossibility to manually screen all of the images on all of the web sites that are added each day to the web. They rely on either storing a local database of website URLs, or referencing the database on the Internet.

Other approach such as described in U.S. patent 5,668,897 by Stoflo (09/16/1997), categorizes images based on a unique image signature into a database for later retrieval and comparison. Such solutions are limited by a known collection of images, which will always be a subset of images created.

Various image-processing algorithms have been investigated for use in detecting objectionable media. For example, algorithms have been tested for use in recognizing shapes, such as people in general, and specific body parts. A detailed summary of work done with algorithms is found in David A. Forsyth and Margaret Flich, Finding Naked People, Journal Reviewing, 1996 and Margaret Flich, David A. Forsyth, Chris Bregler, Finding Naked People, Proceedings of 4th European Conference on Computer Vision, 1996; and David A. Forsyth et al., Finding Pictures of Objects in Large Collections of Images, Proceedings, International Workshop on Object Recognition, Cambridge, 1996. However, all of the above describe individual approaches to analyzing single images using single criteria. None of these publications provide an algorithm/system even close to a robust system, which can be practically used.

Several patents in this field were granted. U.S. patent 6,148,092 to Qian et. al. (11/14/2000) describes a method of detecting skin-tone and in particular detecting faces, using a luminance chrominance algorithm, which is limited to well defined and full bodies. U.S. patent 5,638,136 to Kojima et. al. (06/07/1997) describes yet another method of detecting flesh-tone, and again, this method is limited to well defined chrominance information.

In unrelated fields, Japan patent 09237348A to Hiroshi et. al (09/09/1997) describes a method of determining the posture of a body. Hiroshi et. al. has limited usefulness being again, dependent on color segmentation of an image. US 6,182,081 to Dietl et. al. describes a

method for performing an interactive review of the data contents of a computer with a view to the manual screening of objectionable material contained thereon. However, this method is limited to screening text data against a list of objectionable words and collecting all image data in a thumbnail form for manual review. Thus it is not suitable for application to very large collections of images.

In order for an algorithm to be useful for screening objectionable images, it is necessary for the algorithm to achieve a very high ratio of the number of objectionable images correctly identified to the total number of objectionable images in a database. Unfortunately, no algorithm can determine with full accuracy if an image is of pornographic nature or simply an artistic nude, erotic image or an image with a large amount of skin tone but not of any offensive nature. PCT application of USA application WO00/67204 to Papazian et. al. describes the advantage of using a multiple selection of images to increase the overall likelihood, using the fact that the distribution of the likelihood of detection is spread in a Gaussain fashion and the variance is reduced as a function of the samples. However, Papazian et. al. are not utilizing the cross information that one can achieve from a collection of images, but merely using a statistical improvement.

Similarly in the field of copyright detection. The research work and patents applied all relate to different methods and techniques of watermarking images and then detecting watermarked images. Such techniques are described in EPO EP1/126408 to Wen et. al (22/08/2001) describing a method of detecting embedded information in images. US patent 06,259,801 B1 to Wakasu (07/10/2001) describes watermarking and detecting of watermarked images using DCT methods. US patent publication US2001/0002931 A1 to Maes describes means of detecting images that were marked using geometrical shapes. The drawback in such an approach is that individual detection of watermarked images does not easily or practically lend itself to any form of automatic or workflow solution.

SUMMARY

The primary object of the invention is to provide a more accurate screening tool for suspected pornographic images.

A further object is to increase human body detection accuracy beyond skin-tone criteria.

Another object of the invention is to enable pornographic screening for synthetic images as well as black-and-white or graphic images.

Another object of the invention is to provide faster tools for searching for pornographic images.

A further object of the invention is to provide an enterprise software workflow solution to screen images.

Yet another object of the invention is to improve productivity by combining an automated process for large volume and a secondary manual process for the final decision-making.

Another object of the invention is to provide more efficient tools for searching and detecting copyrighted images.

Briefly, a preferred embodiment of the present invention includes an enterprise based workflow server based system to detect and classify images as potentially of pornographic nature. The system is based on two independent stages. The first stage is an automatic stage which eliminates most images that are positively non pornographic, including an image analysis engine with a plurality of sub engines, each dedicated to different attributes of an image. The second stage describes an additional embodiment of a manual review stage. An alternative embodiment describes a similar server solution to detect images suspected as copyrighted material. The images can be individually classified, or analyzed as a cluster or a collection of images such as an album or motion-picture segment, sampled into key still frames. The system optimizes time saving by elimination of images, and assures accuracy by generating a workflow reviewing solution which can involve a human decision-making stage.

An alternative embodiment describes utilizing the same system with different detection criteria to classify images suspected to be copyrighted.

An advantage of the present invention is that it provides a workflow solution to handle a review of a large amount of digital images.

A further advantage is that the workflow includes an automated step, which eliminates most images and a manual stage, which is of higher accuracy, while dealing with a smaller number of images.

A still further advantage of the present invention is that it provides a modular architecture that is scalable and adjustable based on the load and the load balancing needed.

Another advantage of the present invention is that the automatic classification module can be fine tuned or changed without changing the workflow.

A still further advantage of the method of the present invention is that web sites providing images to the public can block any display of inappropriate pornographic material.

Another advantage of the method of the present invention is that businesses that sell and license images can prevent unlawful usage of such images.

In the Drawing

Fig. 1 is a schematic block diagram of the components in the workflow system;

Fig. 2 shows a flow chart of the workflow;

Figure 3 describes the workflow of an individual analysis engine;

Figure 4 illustrates the manual review workflow;

Figure 4A illustrates the process of setting up priority criteria for the manual reviewing queue;

Figure 5 illustrates the workflow for an alternative embodiment wherein the analysis engine is utilized to detect copyrighted material;

Figure 5A illustrates the components of the imaging engine as related to copyright material;

Figure 5B illustrates the components of the imaging engine as related to offensive graphical material; and

Figure 6 illustrates the user interface look and feel of the reviewer's console.

Detailed Description of the Preferred Embodiment

The method of the present invention applies to the detection of pornographic images transmitted over a communications network, and includes single or still images, motion-picture segments and collections of images such as albums. Alternatively, the system can detect images of copyrighted nature.

In accordance with the present invention, Figure 1 illustrates a schematic diagram of the workflow system components. The system includes end users 300, which subscribe to a photographic community 310, in particular, but not limited to an Internet based photo-sharing site. Images in such community 310 can be stored individually 312, in clusters or otherwise known as collections or albums 314, which can be physical or virtual. Images can be of different file formats and different nature, including and not limited to still images 312 (such as jpeg, GIF, TIF, Jpeg200 etc file formats), or digital movies 316 such as MPEG, AVI, and MP4 etc. A pornography detection server can be on a separate machine connected to the main site via a known Application Protocol Interface (API). The physical connection 320 can

be via a local network, or via the Internet or Intranet. The Porn server 330 also has image detection engines 340 which again can be dedicated CPUs or part of a single configuration.

The engines 340 perform the analysis of the individual images. The engine itself includes scheduler 346 for controlling the retrieval and processing steps for the images, an input stack 344 which retrieves the images in the background, thus saving time and improving the efficiency of the system by eliminating communication and bandwidth delays for retrieval of the images. The engine has sub-engines 350,352,354,356 and 358 each performing a different analysis. The sub-engines include an engine to perform shape analysis of body parts 350, an engine to analyze skin tone 352,an engine to analyze texture of skin 354, an engine to analyze text that accompanies the image, such as title or description 356, and a curvature analysis 358. A likelihood-analyzer 342 receives all information from the sub-engines to create a single value of likelihood also referred to as combined likelihood that the image is pornographic. The statistical likelihood may be expressed as a probability value which is a number expressing the likelihood that the image is of pornographic nature, expressed as the ratio of the number of actual occurrences to the number of possible occurrences. Alternatively, other statistical likelihood measures may be provided based on a variety known of statistical metrics.

All data about the images, including meta-data, which is additional data on the image provided by the site 310, and their classification and categorizations etc. are stored on a central database 360. The images themselves can be saved as a pointer to their storage location on the photo-community 310. Alternatively, the porn server 330 also includes a storage facility for the images 368 which can act as a temporary storage for the images. The images can be in full resolution as provided by the photo community 310. As an additional embodiment, images can be subsampled or reduced in size to an optimal size, which is small enough but still enable good detection of images. The subsampling engine 324 can perform subsampling as part of the transformation of the image from the photo-community 310 to the server 330. In case there are issues of bandwidth, the subsampler module 324 can be located on the photo-community side prior to sending the images to the server. In case there is no bandwidth issues such subsampling can be executed before the images are provided to the engines 340. The image storage component 368 can be for temporary storage of images as part of the review process, or alternatively this storage can be used as the long-term storage for the images as part of the long-term storage of the photo-community, depending on the architecture of the system at large.

The database 360 links to the stored images 368 as well as stores meta data information about the images. When images and meta data are stored in the database in block 360, the database stores them in two tables. The first table 362 is for the individual images, the second 364 is for a collection of images, i.e. albums. Each album points to individual image entries in the image table 362. A third table 366 includes information as to the reviewing process such as a reviewer's log in and password, statistics on each reviewer, priority of sorting images, rating system etc. A queuing system 370 determines the order of images and or albums to be reviewed by the reviewers 380. The reviewers may include a plurality of stations, which are linked to the queuing system 370 via an internal network or even directly through HTML pages on the Internet. Special reviewer 382 may be assign as a supervisor 382 which may determine the priorities, assign reviewers etc.

Referring to the flow diagram of the image analysis workflow in Figure 2, a user 300 uploads a single image or a collection of images 305 to the site 310. The images are stored 307 by the user individually 312 or as a collection 314 or motion pictures 316. Upon receiving a new batch of images, the photo-community server 310 sends a request 315 to the porn server 330. Such a request can be packaged as an XML request or any available communication command between applications. To eliminate redundancies of sending the same images multiple times, and to ensure unique classification of images, each image and each collection should include a unique image identification that both the photo-community 310 and the server 330 can refer to. The Porn-server 330 receives images and places 335 the images and associated information in the database 360. The information is placed in the two tables 362 for individual frames and 364 for collection of frames. In an additional embodiment, the server can also store the images 369 in its storage 368. Alternatively, only a link to the images can be stored in the database, while the images are still linked to their original storage location on the photo-site 310. An additional step can be a subsample of the images 325 before they are stored and processed. Such subsample can occur on the photo community 310 if bandwidth is an issue between the photo-community and the server, or as part of the transfer process 315. In the case of motion picture data, the subsampling process also means selecting key frames forming the movie clip. This is done once again to conserve bandwidth and later on unnecessary repeated calculations by the engine. Individual scenes in a movie are usually a repeat of previous frames with slight modifications, which depicts object movements or camera movements. Therefore, as part of the subsampling process, it is sufficient to only select key frames, one or a few from each scene. Alternatively, a movie collection can be sampled in equal or random frame intervals.

The images are placed in a queue 337 waiting to be processed by the engines 340. An engine, when free, pulls the queue 341 and if images are waiting, retrieves them and process them. A detailed description of the engine 340 functionality is provided in Figure 3. An automatic reply is sent to the server and placed 345 in the associated tables for images and collections of images in the database 360 in the image table 362. Due to the pull nature of this system, many engines can work concurrently on the same database, thus creating a very scalable system. Due to the fact that the engines are in a separate thread, or even on a separate machine, the system can also be highly optimized for stress and load balance.

The automatic reply can be made available to the photo-site using the same protocol 320 that was utilized to place the images in the queue. Such a reply can be individual likelihood results for an image, or a collective likelihood for an album. The reply can be sent to the photo-community 339 as a job is finished. Alternatively, the photo-community can query regarding a specific image, or a collection as defined in step 319. Such a reply or query can include a list of most suspected images, or a sorted list by likelihood, wherein the most suspected images are provided up front. Alternatively, the results can be queried based on unique image identification.

Referring to the flow diagram of the engine 340 as illustrated by Figure 3, after being placed in the database, images and or collections of images are placed in a queue. It is the role of the engines, which can be one or many, to query the queue. The management and scheduling of the different analysis an image goes through is all controlled by the scheduler 346. The first task is to retrieve 402 the image and associated meta data from the server, and specifically from the database table 362 and the image storage 368. The Input process is separate from the actual analysis to allow multi processing. While one job is being analyzed, the next job is being retrieved 402 and placed 406 in a temporary storage of the engine. By doing so, the engine can be optimized for an optimal tradeoff between bandwidth to get the images and time to analyze them.

In the case of a movie clip, an additional step may occur wherein the movie clips are segmented and still frames are extracted, as described in step 325. The retrieved images are then sent by the scheduler to the individual analysis modules. Images are sent in step 450 to the shape analysis engine to detect indicative shapes associated with body parts. Independently, images are sent in step 452 to the skin tone analysis engine where human bodies are classified based on color attributes. Similarly, images are sent 454 to the Texture analysis engine where the picture is analyzed to identify bodies based on texture characteristics. Any textual meta data associated with an image and or collection, such as

image names, comments etc. are sent 456 to the textual analysis engine, where suspicious words and sentences are being patterned and matched. The images are also sent 458 to a curvature analysis sub-engine 358 for possible matching with shapes based on curves.

Alternatively, each or some of the sub-engines may be concatenated. For example, the skin-tone sub engine may classify some regions. This information along with the image will be then sent to the curvature engine and the latter will examine only images marked by the former engine. Any combination and ordering may be executed by the scheduler to improve the accuracy and speed of the process. Additionally, other sub engines can be added to the process if deemed necessary. Such engines may be based on specific meta data that is provided by the site 310, or any improved detection technology. Such additions of new sub engines can be done transparently to the process workflow process and therefore the advantage of the system as defined.

The data from all sub engines is then analyzed 420 by the combined likelihood calculator 342. A single parameter is then generated to describe the overall probability that an image is of pornographic nature. The formulae for this calculation can be tweaked and optimized as the process matures. The individual probabilities along with the combined likelihood are then returned to the main database.

Fig. 4 shows an additional embodiment wherein images can be sent to a manual reviewing engine for a definitive classification. As a preparatory step, supervisor 382 can determine 482 some rules 385 as to the order that images should be reviewed. Figure 4A. illustrates such priority list which can be modified by the supervisor. The supervisor may assign primary, secondary tertiary etc. criteria for sorting the queue based on different criteria. Such rules can be to review images based on the statistic likelihood of individual images 800, the collective likelihood of albums 802, percentage of images suspected 804, or the total number of images suspected 806 etc.

Other criteria can include additional information or meta data associated with the images such as the number of pageviews 820, which is the number of times an image was viewed, sorting by the date the image was uploaded 810, or the time the image was in the queue 812.

Other criteria may include meta data provided by the photo-community 310 such as prior offenses that a user may have 830, etc. In figure 4A as an illustrative example, the supervisor marked the first sorting criteria 861, to be pageviews 830, second sorting 862 to be the statistical likelihood 810 and the third criteria 863 to be the upload date and time 802.

Reviewers, which can be one or many, can login 486 and logout 512. While 488 images are still in the review queue 370, any free reviewing station pulls the next image or selection of images within in the queue. An album is being displayed 492 on the reviewer's screen as further illustrated in Figure 5. The interface between the reviewing station and the database can be via some secure HTML communication or any Interface language such as XML. In such a manner, reviewers may log into the system remotely and out of the site's internal security system such as a firewall. With this architecture, multiple reviewers can be added as needed based on the load of the system and the productivity of the reviewers.

The reviewers can assign 494 classification of individual images or a collection. If a reviewer is not sure as to the rating of an image or a collection as defined in step 496, the images can be sent 498 for a second and deciding opinion by the supervisor. The supervisor will then repeat similar reviewing process as the reviewer did.

Upon assigning a rating, either by the reviewer or a supervisor, the data is returned 510 to the database 360. The results from the database can be sent to the photo community or asynchronously retrieved by the photo community upon request as described in Figure 2 blocks 339 and 319 respectively.

Figure 6 depicts an example of a potential user interface of the reviewer. The window in which a reviewer works can be HTML based. The information included includes an action window 610 wherein the reviewer can toggle between views, go to the next album or send albums to the supervisor when the reviewer is not sure as described in Figure 4 block 498. The information also includes a text window 600 with the meta data that may help the reviewer in deciding as to the rating, such as the number of hits, which is also referred to as "page views", or otherwise number of people that looked at the images, additional text which alternatively can be highlighted or marked by the textual analysis engine 356, the ration between the total number and the suspected number of images etc. In terms of the user interface, any parameter that is beyond the acceptable normal can be marked differently, such as highlighted, marked in red etc. to immediately attract the reviewer's attention. Object 640 displays the pertinent information as to the current reviewed album. Block 630 displays the images themselves. The images may be displayed in reduced size, also termed as thumbnails, or as full resolution. The reviewer may decide to view all images of the album or toggle to an alternative view where only the suspected images are displayed. This workflow will enable fast detection of suspected albums. The reviewer can then assign a rating, which can be redefined by the site. Different sites may have different needs, which vary by countries, and audience. One rating which is very important is the ability to mark images as illegal as

illustrated in block 620. Illegal images include for example child pornography. Such albums are immediately moved up the queue and may be reported especially to the server 330 and to the client photo-community 310 for immediate action.

As alternative embodiments, due to the modular structure of the server, additional components can be easily added if needed. For example the detection engines may extend beyond pornographic detection and include a detection engine for copyrighted material. This is a useful tool to block images that have been unlawfully taken from other sources such as other webs sites and other publications. The reviewing system is identical to the one as explained in the main embodiments. All the difference is in the type of sub-engines used within the engine 340. Figure 5A illustrates the components of the engine 740 as related to copyright material. Such a tool can be of great use for companies that are trying to protect their visual assets. Examples include celebrities, news agencies and image banks.

Referring to components of the imaging engine 740, which is the engine 340 modified to support copyright detection as illustrated by Figure 5A, the sub-engines that will be used include a watermarking detection engine 751, a halftone pattern recognition engine 753, a textual analysis engine 755 and a graphical analysis engine 757. The textual analysis may also detect unique signatures that an image may have in its header or meta data which can link it to its original owner. The graphical analysis engine may detect known copyrighted or trademarked logos, signs or figures.

Referring to workflow of the copyright violation detection system illustrated by Figure 5, the images will pass through the sub engines described in Figure 5A to detect if images include watermarking 750 in them; detect the possibility that images were unlawfully scanned from printed publications such as magazines 752, search for key words such as famous people, current events news etc. 754, and search for known graphical logos and symbols 756.

Of course, a hybrid system that can include both pornographic material and copyright violation material can be constructed by combining the pornographic detection engines 350, 352, 354, 356, and 358 with the copyright violation detection 751, 753 and 755.

As further embodiments of the present invention, due to the modular structure of the server, additional detection engines may extend beyond pornographic detection and copyright violation detection to include a detection engine for graphically-offensive material. This is a useful tool to prevent for example hate-crime related material or offensive political propaganda material. As an example, such material may include the presence of the Swastika symbol in images, logos of supremacy organization etc. As described before, the reviewing

system is identical to the one as explained in the main embodiments. All the difference is in the type of sub-engines used within the engine 340. Figure 5B illustrates the components of the engine 840 as related to offensive graphical material.

Graphical images can be stored digitally in various ways. In particular, they can be of photographic pixel representation, which is a grid based representation of image, artificial line-art representation such as drawings of art objects and vector representation which is a mathematical formulae to represent separate shapes such as lines, polygons and text, and groups of such objects, as opposed to bitmaps.

Referring to components of the imaging engine 840, which is the engine 340 modified to support detection of offensive material as illustrated by Figure 5B, the sub-engines that will be used include a line-art analyzer 790, a pixel analyzer 792 and a vector analyzer 794.

In conclusion, the reader may see that the workflow solution as described in the present invention can be used to increase productivity and improve the detection capability of imaging based reviewing systems. Specifically, such a system can ensure the absence of pornographic material on community based web sites.

In addition to the increased throughput of the automated system, the present invention also describes a manual workflow stage in which reviewers can efficiently review images and classify them. Furthermore, due to the modular structure of this architecture, new detection components can be added without the need to redesign or affect the workflow and or the interface of this system. Furthermore, such server-based image reviewing solution can be utilized for other purposes involving inappropriate use of images, such as the detection of copyrighted material.

Although the description above contains many specificities, these should not be construed as limiting the scope of the invention but as merely providing illustrations of some of the presently preferred embodiments of this invention. Various other embodiments and ramifications will be apparent and possible within it's scope for those skilled in the art. Thus the scope of the invention should be determined by the appended claims and their legal equivalents as covering all such alterations and modifications that fall within the true scope and spirit of the invention, rather than by the examples given.

In the claims: